



14 MAR 2005 - - - - -



INVESTOR IN PEOPLE

Rec'd PCT/PTO 14 MAR 2005

The Patent Office

Concept House

Cardiff Road

Newport REC'D 02 DEC 2003
South Wales

NP10 8QQ WIPO

PC1

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated 17 November 2003

PRIORITY DOCUMENT

SUBMITTED OR TRANSMITTED IN
COMPLIANCE WITH RULE 17.1(a) OR (b)

13SEP02 E747348-1 DO2000
P01/7700 0.00-0221144.91/77
The Patent OfficeCardiff Road
Newport
South Wales
NP10 8QQ**Request for grant of a patent**

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

1. Your reference

PDG/24508

2. Patent application number

(The Patent Office will fill in this part)

0221144.9

12 SEP 2002

3. Full name, address and postcode of the or of each applicant (underline all surnames)

Snell & Wilcox Limited
Principal Engineer - Patents & IP
Durford Mill
Petersfield
Hampshire GU31 5AZ 5579784004
United Kingdom

Patents ADP number (if you know it)

If the applicant is a corporate body, give the country/state of its incorporation

4. Title of the invention

Image Processing Using Vectors

5. Name of your agent (if you have one)

"Address for service" in the United Kingdom to which all correspondence should be sent (including the postcode)

MATHYS & SQUIRE
100 Gray's Inn Road
London WC1X 8AL
United Kingdom

Patents ADP number (if you know it)

1081001

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number

Country

Priority application number
(if you know it)Date of filing
(day / month / year)

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application

Date of filing
(day / month / year)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:

YES

- a) any applicant named in part 3 is not an inventor, or
- b) there is an inventor who is not named as an applicant, or
- c) any named applicant is a corporate body.

See note (d)

9. Enter the number of sheets for any of the following items you are filing with this form.
Do not count copies of the same document

Continuation sheets of this form
Description 7
Claim(s) —
Abstract — *(Handwritten mark)*
Drawing(s) —

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (Patents Form 7/77)

Request for preliminary examination and search (Patents Form 9/77)

1

Request for substantive examination (Patents Form 10/77)

Any other documents
(please specify)

11.

I/We request the grant of a patent on the basis of this application.

Signature

Mathy S
MATHYS & SQUIRE

Date

12 September 2002

12. Name and daytime telephone number of person to contact in the United Kingdom

GARRATT, Peter Douglas - 020 7830 0000

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 08459 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.

IMAGE PROCESSING USING VECTORS

This invention is directed to a method of image segmentation using a vector quantization approach.

- 5 The aim of the algorithm is to partition the picture into a fixed number of segments, taking into account the location and value of each pixel. For moving sequences, we additionally desire a smooth transition in the segmentation from one picture to the next.
- 10 We start by representing the picture in multidimensional space, which we shall call the universe. This is the product of the ordinary two-dimensional space of the picture, which we shall call the canvas, and the pixel space, which is the space in which the pixel values themselves are defined. For example, if we are segmenting an RGB picture, the pixel space will have three dimensions and the
- 15 universe will therefore have five dimensions.

The picture is represented as a set of points in the universe, one for each pixel. The co-ordinates of the pixel in the universe describe its position on the canvas together with its value in pixel space.

- 20 The segmentation is simply a partitioning of the set of pixels into a fixed number of subsets.

- 25 The algorithm works as follows. We start with some initial segmentation of the picture into the desired number of segments. At the very start of a sequence, this might correspond to a straightforward division of the canvas into equal areas. Subsequently, the initial segmentation may be the segmentation of the previous picture. We then perform two steps.

First, we calculate the centroid of each segment, according to the values of the segment's pixels in the universe.

We then reassign each pixel to the segment with the closest centroid, where

5 closeness is measured by the Euclidean distance in the universe.

These two steps may be repeated once or more using the same picture. The result is a new segmentation, which may be used as the initial segmentation for the next picture.

10

A segmented version of the picture can be created by replacing the pixel values in each segment with the projections of the segment centroids onto the pixel space.

15 The algorithm has the following features, each of which may be optionally combined with any of the others:

§ all segments have the same mass.

§ the segmentation operates recursively, starting from an initial

20 segmentation

§ pixels are assigned to clusters and then the clusters are moved to reflect their new membership

The algorithm performs more consistently than previous approaches

25

The algorithm can actually be considered as an adaptive vector quantization process. The universe is the vector space, and the codebook is the set of segment centroids. The codebook is updated according to the data being coded.

Vector quantization helps to explain the algorithm as a compression technique in which the goal is to describe the information using a limited set of representative points in the universe.

5 We now turn to some details of the algorithm.

The performance of the algorithm depends on the relative scaling of the co-ordinate axes in the universe. For example, if on the one hand the canvas co-ordinates are multiplied by a large scaling factor, the algorithm will tend to pay 10 more attention to the spatial component of the overall Euclidean distance and the final segmentation will look more like a simple partitioning of the canvas. If on the other hand the pixel space co-ordinates are given a large scaling factor, the segmentation will be dominated by pixel values and each segment will be spread across the canvas, giving a result closer to simple histogram analysis.

15

Good results have been obtained by setting the relative scaling factors of the co-ordinate axes so as to equalize the variances of the pixels in all the dimensions of the universe. An exception to this rule is that the canvas co-ordinates are arranged to have equal scaling determined by the wider (usually horizontal) co-ordinate. Setting the scaling according to variance is equivalent to using the so-called Mahalanobis Distance, as described for example in
http://www.engr.sjsu.edu/~knapp/HCIRODPR/PR_Mahal/M_metric.htm.
in the special case where the pixel co-ordinates in the universe are assumed to be uncorrelated.

25

Other embodiments employ methods of dynamically varying the relative scaling in order to minimize some error measure. This is difficult, as the obvious error measures available are directly dependent on the scaling. In one embodiment, the overall vector quantization is expressed as the product of the errors in each 30 co-ordinate, with the constraint that all the scaling factors sum to a constant.

In the algorithm, the number of segments is a parameter chosen by the user. In alternatives, the number of segments is chosen as a function of the input data. For example, the number of segments may be chosen so that the variance of the overall vector quantization error approaches a predetermined constant value.

5

It is possible that a segment may disappear after some frames if its centroid turns out to be further from every pixel than some other segment. In many cases, this is desirable behaviour, as when an object disappears from the screen. In other embodiments, the mechanism may allow for the introduction of new segments.

10 Particular algorithms may establish criteria for splitting an existing segment into two, and possibly also for merging segments that end up being close together in the universe. The criteria could be based on the variance of the vector quantization error, or on the kurtosis of the distribution of the pixels in each segment as a measure of bimodality.

15

Another approach to deciding whether to add or remove segments is to run two or more parallel versions of the algorithm with different numbers of segments and to base the decision on the difference in overall error between the two versions.

20 One possible solution to the problem of the disappearance and reappearance of objects because of global motion in the scene is to impose a toroidal structure on the canvas. This is done by stitching the left edge to the right edge and the top edge to the bottom edge. Centroids that disappear off one edge will now reappear at the opposite edge and will be available for re-use. Care needs to be
25 taken in the distance definition to make sure that the shortest distance is used.

Other variations on the basic Euclidean distance are also possible. One possibility, which is simpler to implement in hardware, is the Manhattan distance or L-1 norm. This is the distance between two points measured by walking
30 parallel to the co-ordinate axes. Another possibility is to take the maximum of the differences measured along the co-ordinate axes.

Certain embodiments use a pixel space in which the co-ordinates are simply the luminance, YUV or RGB values of the pixels. Others look at segmentation on the 5 basis of motion vectors and this is described in a separate section below.

There are many other possibilities for features that can be included in the pixel space, for example local measures of noise or texture, or (in the case of the Delft 10 presentation referred to above) Gabor jets. For this reason, the pixel space (or the whole universe) is sometimes referred to as the feature space.

The description so far assumes that each pixel belongs to just one segment, so that assignment of pixels to segments is a hard decision. It is possible to replace this with a soft decision, in which each pixel carries a set of probabilities of 15 membership of each segment. The output of the algorithm can be based on a hardened version of the decision (assigning the segment with the highest probability) while the soft decision is retained for the recursive processing. Alternatively, the soft decision may have significance in the output of the algorithm. For example, if the pixel space consists of motion vectors, the output 20 may consist of several motion vectors assigned to each pixel, each with a relative weight.

The process of calculating centroids and assigning vectors can be repeated several times per picture. Performing several iterations is necessary if the aim is 25 to segment a single picture starting from trivial initial conditions. However, I have found that for moving sequences there may be little to be gained from performing more than one iteration per picture.

The following describes additional features of the algorithm in the case where the 30 pixel space includes motion or displacement vectors.

The algorithm can be used to segment a motion vector field resulting from another process such as Phase Correlation (PhC). Segmentation of the motion vector field fits very well with a model in which a picture consists of distinct objects, each with its own speed and direction of motion. However, it does not

- 5 work so well if the objects are rotating, receding or approaching or if their distance from the camera is not uniform. A better description of the motion in each object can be obtained by replacing the basic motion vector with an affine transform, in which the horizontal and vertical motion vector components are each allowed to vary linearly with the spatial co-ordinates of the pixel on the
- 10 canvas.

The affine motion model can be applied to the segmentation algorithm by expressing each segment centroid as a set of six parameters describing an affine transform. The pixel-space component of the distance metric then becomes the

- 15 distance between the motion vector at the point concerned and the affine model evaluated at that point.

- 20 Motion vectors give rise to the possibility of an additional component in the distance metric, based on the error in the pixel domain when the motion vector is applied to a picture. This displaced frame difference can be incorporated into the distance metric with appropriate scaling. The result is a segmentation of the motion vector field that takes into account the fidelity of the motion compensated prediction.

- 25 Some features of the performance of the algorithm:

§ the variation in the segmentation from one frame to the next is small. This is difficult to achieve with segmentation algorithms that do not use recursion

- 30 § with the scaling chosen, the segments are not necessarily contiguous but the parts of each segment remain close together.

Simulations show that the property of smooth variation is quite reliable and that the segmentation is good provided a suitable number of segments is chosen. This "suitable number" may be chosen as sequence dependent.

- 5 Replacement of "raw" PhC motion vectors by a segmented version can improve the overall quality of the motion vector field. This is true on occasions of the version based on affine transforms.

10 The algorithm performs fast and (subjectively) very well on demanding picture material.

It will be appreciated by those skilled in the art that the invention has been described by way of example only, and that a wide variety of alternative approaches may be adopted.